arXiv:quant-ph/0512082v1 12 Dec 2005

Qubit Complexity of Continuous Problems*

A. Papageorgiou† and J. F. Traub‡
Department of Computer Science, Columbia University, New York, USA
(Dated: February 1, 2008)

The number of qubits used by a quantum algorithm will be a crucial computational resource for the foreseeable future. We show how to obtain the classical query complexity for continuous problems. We then establish a simple formula for a lower bound on the qubit complexity in terms of the classical query complexity.

PACS numbers: 03.67.Lx, 02.60-x
Keywords: Complexity, numerical approximation, quantum algorithms

I. INTRODUCTION

There are two major motivations for studying algorithms and complexity of continuous problems.

1. Many scientific problems have continuous formulations. Examples include path integration, Feynman-Kac path integration, and the Schrödinger equation.

2. Are quantum computers more powerful than classical computers for important scientific problems? How much more powerful?

To answer these questions one must know the classical computational complexity of the problem. There are especially constructed problems such as Simon's problem [1] for which the quantum speedup is known; see also [2]. Furthermore, it is known that quantum computers enjoy quadratic speedup for search of an unordered database [3]. Knowing the quantum speedup for such a discrete problem is the exception. Generally, for discrete problems we do not know the computational complexity. (Examples of discrete problems are 3-SAT and the traveling salesman problem.) We have to settle for the conjecture that the complexity hierarchy does not collapse. A famous example of this conjecture is that P $\neq$ NP. Thus although is is widely believed that Shor's algorithm [4] gives an exponential speedup it is only a conjecture because the classical computational complexity of integer factorization is an important open problem.

In what follows it is important to stress the difference between the cost of an algorithm for solving a given problem, and the computational complexity of this problem. The computational complexity (for brevity, the complexity) is the *minimal* computational resources needed to solve the problem. Examples of computational resources, which have been studied, include memory, time, and communication on a classical computer and qubits, quantum gates and queries on a quantum computer. For the foreseeable future qubits will be a limiting resource and in this paper we'll give a general lower bound on the qubit complexity for continuous problems.

For continuous problems we often know the classical complexity. There is a large literature in the field of information-based complexity which studies problems with partial and/or contaminated information; see [5, 6] and the references therein. Since functions of a continuous variable cannot generally be input into a digital computer, the computer has only partial information about them. As we shall see in Section II this makes it possible to use an adversary argument to get a lower bound on the classical query complexity, and, therefore, on the total complexity of many continuous problems.

Most continuous problems arising in practice cannot be solved analytically; they must be solved numerically. Since a digital computer has only partial information about the input function the problem can only be solved approximately, to within an error threshold ε. If one insists on an error at most ε for all inputs in a class F (the worst case setting) it's been shown that for many multivariate problems the complexity is exponential in the number of variables. This is known as the curse of dimensionality and such problems are said to be intractable. Note that for continuous problems many problems are known to be intractable while for discrete problems the intractability of NP-hard problems is only conjectured; see Remark II.4.

There are two major ways to break the curse of dimensionality, see [6, p. 24]. We can weaken the worst case assurance, accepting instead a stochastic assurance such as in the randomized setting. The Monte Carlo algorithm is known to be optimal for integration in this setting if F is the class of bounded

*This research was supported in part by DARPA and NSF.
†Electronic address: ap@cs.columbia.edu
‡Electronic address: traub@cs.columbia.edu

continuous functions. Or we can change the class F of inputs. By suitable choices of F we can sometimes provide a worst case guarantee while breaking intractability.

We outline the remainder of the paper. In Section II we illustrate the adversary argument which will provide us with the classical information complexity. We use a very simple example to do this. In Section III we provide a more general formulation and introduce notation. In the concluding section we'll prove a general theorem giving a lower bound on the qubit complexity in terms of the classical query complexity.

II. CLASSICAL INFORMATION COMPLEXITY

We will illustrate the adversary argument used to obtain the classical query complexity using a very simple example. The same idea can be applied very generally [5, 6].

We want to compute $I(f) = \int_0^1 f(x)\,dx$. Call f the mathematical input. For most integrands we can't use the fundamental theorem of calculus to compute the integral analytically; we have to approximate it numerically. Although we can input the symbolic form of f into a digital computer it doesn't help us to compute the integral. We compute

$$y_f = [f(t_1), \ldots, f(t_n)]$$

at n a priori chosen deterministic points t_i, $i = 1, \ldots, n$. Given y_f, there are an infinite number of functions with the same y_f. That is, we have only partial information about the mathematical input. Even though the functions may have the same y_f their integrals may be very different. Let G be the set of functions with the same y_f; see Figure 1.

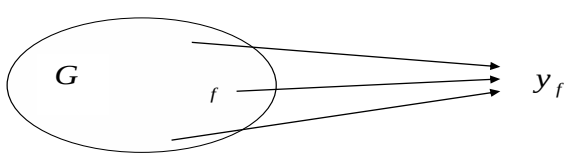

FIG. 1: Indistinguishable functions

If we only assume that f is, say, Riemann-integrable the classical query complexity is infinite, i.e, we cannot achieve any desired accuracy no matter how large n is. To get finite complexity we have to make a promise about f. With the promise that the absolute value of the functions under consideration is uniformly bounded by a known constant is it is easy to show the complexity is still infinite. Thus we further restrict the class of inputs and assume that our function belongs to

$$F = \{f : \ |f'(x)| \le L, \quad x \in [0,1]\}.$$

Let $K = F \cap G$. The functions in K are indistinguishable; see Figure 2.

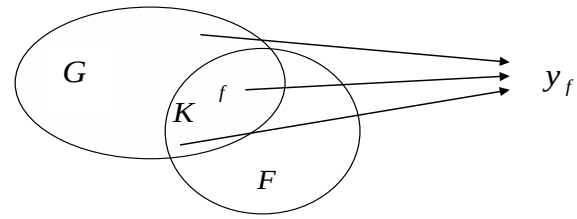

FIG. 2: Restrict the class of functions

Let H denote the set $I(\tilde{f})$, $\tilde{f} \in K$. It is easy to show that H is an interval and that its length varies with y_f, n and the points $t_1, \ldots, t_n$. Any number in H is a potential approximation to the integral. A measure of the intrinsic uncertainty in our approximations is the size of H. There is a standard concept of the size of a set; it is the radius of the smallest ball containing the set. We call this radius the radius of information, rad, because its magnitude depends on how much information we have about the true f. It is easy to show that we can guarantee an ε-approximation iff $rad \le \varepsilon$. Let $m(\varepsilon)$ be the minimum number of function evaluations needed to solve the problem to within ε. The condition $rad \le \varepsilon$ implies that if we compute less than $m(\varepsilon)$ function evaluations there does not exist any algorithm which solves the problem with error ε. See [6, Section II.2] for a general discussion of the radius of information.

Let $\mathbf{c}$ be the cost of a query, that is of a function evaluation. We define the classical query complexity, $\mathrm{comp}^{\mathrm{query}}_{\mathrm{clas}}(\varepsilon)$, as

$$\mathrm{comp}^{\mathrm{query}}_{\mathrm{clas}}(\varepsilon) = \mathbf{c}\, m(\varepsilon). \tag{1}$$

The query complexity is the minimum amount that must be paid to obtain the information about f needed to compute $I(f)$ to within ε.

In the concluding section we will see how the classical query complexity is used to lower bound the quantum qubit complexity. We conclude this section with some remarks.

Remark II.1. *The type of argument we have used in this section is called an adversary argument because if we don't collect enough information an imagined adversary can claim the mathematical input is a function g for which $I(g)$ is very different than $I(f)$, foiling the assurance that we've computed an ε-approximation to $I(f)$.*

Remark II.2. *Note that there has been no mention of how $y_f = [f(t_1), \ldots, f(t_n)]$ is used to approximate the integral $I(f)$. This can be done by an algorithm ϕ of the form*

$$\phi(f) = \sum_{j=1}^{n} a_j f(t_j).$$

There is a large literature on the optimal choice of the coefficients a_j in ϕ, and on the optimal choice of the points t_j; see, for example [6]. Part of the power of the approach we've illustrated here is that decisions concerning information can be separated from decisions regarding algorithms.

Remark II.3. *The mathematical tools for lower and upper bounds on classical query complexity (and other types of complexity) are often deep but this example gives the idea of the adversary argument.*

Remark II.4. *Why can we obtain the complexity of continuous problems whereas we have to settle for conjectures about the complexity hierarchy for discrete problems? For continuous problems we have partial information and we can use the adversary argument to get lower bounds. For discrete problems we have complete information. For example, for the traveling salesman problem we are given the locations of the cities and these coordinates can be input into a digital computer. There is no information level and no adversary argument.*

III. FUNDAMENTAL CONCEPTS AND NOTATION FOR QUANTUM COMPUTATION

A quantum algorithm consists of a sequence of unitary transformations applied to an initial state. The result of the algorithm is obtained by measuring its final state. The quantum model of computation is discussed in detail in [2, 7, 8, 9, 10, 11]. We summarize this model to the extent necessary for this paper.

The initial state $|\psi_0\rangle$ of the algorithm is a unit vector of the Hilbert space $\mathcal{H}_\nu = \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2$, ν times, for some appropriately chosen integer ν, where $\mathbb{C}^2$ is the two dimensional space of complex numbers. The dimension of $\mathcal{H}_\nu$ is 2^ν. The number ν denotes the number of qubits used by the quantum algorithm.

The final state $|\psi\rangle$ is also a unit vector of $\mathcal{H}_\nu$ and is obtained from the initial state $|\psi_0\rangle$ through a sequence of unitary $2^\nu \times 2^\nu$ matrices, i.e.,

$$|\psi\rangle := U_T Q_f U_{T-1} Q_f \cdots U_1 Q_f U_0 |\psi_0\rangle. \qquad (2)$$

The unitary matrix Q_f is called a quantum query and is used to provide information to the algorithm about a function f. Q_f depends on n function evaluations $f(t_1), \ldots, f(t_n)$, $n \le 2^\nu$. The $U_0, U_1, \ldots, U_T$ are unitary matrices that do not depend on f. The integer T denotes the number of quantum queries.

For algorithms solving discrete problems, such as Grover's algorithm for the search of an unordered database [3], the input f is considered to be a Boolean function. However, classical algorithms solving continuous problems using floating or fixed point arithmetic can also be written in the form of (2). Indeed, all classical bit operations can be simulated by quantum computations, see e.g., [8].

The most commonly studied quantum query is the *bit* query. For a Boolean function $f : \{0, \ldots, 2^m - 1\} \to \{0, 1\}$, the bit query is defined by

$$Q_f |j\rangle |k\rangle = |j\rangle |k \oplus f(j)\rangle.$$

Here $\nu = m + 1$, $|j\rangle \in \mathcal{H}_m$, and $|k\rangle \in \mathcal{H}_1$ with $\oplus$ denoting the addition modulo 2. For a real function f the query is constructed by taking the most significant bits of the function f evaluated at some points t_j. More precisely, as in [10], the bit query for f has the form

$$Q_f |j\rangle |k\rangle = |j\rangle |k \oplus \beta(f(\tau(j)))\rangle,$$

where the number of qubits is now $\nu = m' + m''$ and $|j\rangle \in \mathcal{H}_{m'}$, $|k\rangle \in \mathcal{H}_{m''}$. The functions β and τ are used to discretize the domain $\mathcal{D}$ and the range $\mathcal{R}$ of f, respectively. Therefore, $\beta : \mathcal{R} \to \{0, 1, \ldots, 2^{m''} - 1\}$ and $\tau : \{0, 1, \ldots, 2^{m'} - 1\} \to \mathcal{D}$. Hence, we compute f at $t_j = \tau(j)$ and then take the m'' most significant bits of $f(t_j)$ by $\beta(f(t_j))$, for the details and the possible use of ancillary qubits see [10].

At the end of the quantum algorithm, a measurement is applied to its final state $|\psi\rangle$. The measurement produces one of M outcomes, where $M \le 2^\nu$. Outcome $j \in \{0, 1, \ldots, M - 1\}$ occurs with probability $p_f(j)$, which depends on j and the input f. Knowing the outcome j, we compute classically the final result $\phi_f(j)$ of the algorithm.

In principle, quantum algorithms may have many measurements applied between sequences of unitary transformations of the form presented above. However, any algorithm with many measurements can be simulated by a quantum algorithm with only one measurement at the end [8].

We are interested in continuous problems such as multivariate and path integration, multivariate approximation, ordinary and partial differential equations, and the Sturm-Liouville eigenvalue problem. For many continuous problems we know tight quantum complexity bounds [10, 12, 13, 14, 15, 16, 17, 18].

Let S be a linear or nonlinear operator such that

$$S : \mathcal{F} \to \mathcal{G}. \tag{3}$$

Typically, $\mathcal{F}$ is a linear space of continuous real functions of several variables, and $\mathcal{G}$ is a normed linear space. We wish to approximate $S(f)$ to within ε for $f \in \mathcal{F}$. We approximate $S(f)$ using n function evaluations $f(t_1), \ldots, f(t_n)$ at deterministically and a priori chosen sample points. The quantum query Q_f encodes this information, and the quantum algorithm obtains this information from Q_f.

Without loss of generality, we consider algorithms that approximate $S(f)$ with probability $p \geq \frac{3}{4}$. The local error of the quantum algorithm (2) that computes the approximation $\phi_f(j)$, for $f \in \mathcal{F}$ and the outcome $j \in \{0, 1, \ldots, M-1\}$, is defined by

$$e(\phi_f) \;=\; \min\left\{ \alpha : \sum_{j:\ \|S(f)-\phi_f(j)\| \leq \alpha} p_f(j) \geq \tfrac{3}{4} \right\}, \tag{4}$$

where $p_f(j)$ denotes the probability of obtaining outcome j for the function f. The *worst probabilistic* error of a quantum algorithm ϕ is defined by

$$e^{\text{quant}}(\phi) \;=\; \sup_{f \in \mathcal{F}} e(\phi_f). \tag{5}$$

IV. LOWER BOUND ON QUBIT COMPLEXITY

For the foreseeable future the number of qubits used by a quantum algorithm will be a crucial computational resource. We will show how to obtain a lower bound for the number of qubits needed for algorithms that approximate continuous problems such as (3). In particular, let $\text{comp}^{\text{qubit}}(\varepsilon)$ be the minimal number of qubits required by a quantum algorithm of the form (2) approximating $S(f)$ with accuracy ε and probability at least $\frac{3}{4}$.

We will derive a lower bound for the qubit complexity using facts about the classical complexity of continuous problems. A similar lower bound result was announced by H. Woźniakowski at the DARPA PI meeting in Chicago in May 2004; see [19] for his proof. The proof we present here is different and constructive. In the analysis of classical algorithms one considers the classical query cost, which depends on the number of function evaluations n used by the classical algorithm. It suffices to consider deterministic classical algorithms ϕ in the worst case, i.e., to measure the error by

$$e^{\text{wor}}(\phi, n) = \sup_{f \in F} \big\| S(f) - \phi(f(t_1), \ldots, f(t_n)) \big\|. \tag{6}$$

The classical query complexity, $\text{comp}^{\text{query}}_{\text{clas}}(\varepsilon)$, of the problem (3) is the minimal number of function evaluations that are necessary for accuracy ε times the cost of a query, i.e.,

$$m(\varepsilon) = \min\left\{ n :\ \exists\ \phi \text{ with } e^{\text{wor}}(\phi, n) \leq \varepsilon \right\},$$

$$\text{comp}^{\text{query}}_{\text{clas}}(\varepsilon) = \mathbf{c}\ m(\varepsilon). \tag{7}$$

The classical query and combinatorial complexities of many continuous problems are known [5, 6]. We are now ready to show how to use classical query complexity lower bounds to derive qubit complexity lower bounds.

Recall that quantum algorithms may require some classical computations to be performed, for instance, at the end after the measurement to produce the final result, or at the beginning to prepare the initial state. These classical computations may or may not include a number of function evaluations. To exclude trivial cases that reduce the qubit complexity at the expense of classical computations, we will assume that the number of function evaluations computed by the classical components of the quantum algorithm cannot exceed the number of function evaluations obtained in superposition by the query due to quantum parallelism.

Theorem IV.1. *The qubit complexity of a quantum algorithm (2) that solves the problem (3) with accuracy ε is bounded from below as follows*

$$\text{comp}^{\text{qubit}}(\varepsilon) \geq \log_2 \left[\text{comp}^{\text{query}}_{\text{clas}}(3\varepsilon)\right] - 1.$$

Proof: Consider a quantum algorithm that solves the problem with accuracy ε. This algorithm uses Q_f which, in turn, depends on a number of function evaluations of f which we denote by $n(\varepsilon)$. It follows that the number of qubits of the quantum algorithm is at least $\log_2 n(\varepsilon)$.

A quantum algorithm that approximates (3) with accuracy ε can be simulated by a classical algorithm. The computational cost of this simulation is not important here. The important fact is that the classical algorithm also uses $n(\varepsilon)$ function evaluations and approximates $S(f)$ with worst probabilistic error (5) less than ε.

Since the algorithm achieves error ε, the final state of the quantum algorithm, and the corresponding state of its classical simulation, contain outcomes j such that $\|S(f) - \phi_f(j)\| \leq \varepsilon$, where the sum of their probabilities is $\sum_j p_f(j) \geq \frac{3}{4}$. Moreover, the classical simulation can compute the probabilities of all the outcomes, since it has computed all the amplitudes in the final state of the quantum algorithm.

The quantities $p_f(j)$ and $\phi_f(j)$, for all possible outcomes j, suffice for computing deterministically an approximation of $S(f)$ with error 3ε. To see this observe that the local error (4) of a quantum algorithm can be equivalently rewritten as

$$e(\phi_f) = \min_{A:\,\mu(A)\geq\frac{3}{4}} \max_{j\in A} \big\|S(f)-\phi_f(j)\big\|, \qquad (8)$$

where $A \subset \{0,1,\ldots,M-1\}$ and $\mu(A) = \sum_{j\in A} p_f(j)$. Consider all sets of outcomes where the sum of the respective probabilities is at least $\frac{3}{4}$. From these discard any set that contains outcomes $j \neq k$ such that $\|\phi_f(j)-\phi_f(k)\| > 2\varepsilon$.

Let A denote one of the remaining sets of outcomes then $\mu(A^*) = \sum_{j\in A^*} p_f(j) \geq \frac{3}{4}$ and $\|\phi_f(j) - \phi_f(k)\| \leq 2\varepsilon$, $j,k \in A$. The fact that $e(\phi_f) \leq \varepsilon$, equation (8) and the triangle inequality imply that A exists.

There exists $j^* \in A$ such that $\|S(f)-\phi_f(j^*)\| \leq \varepsilon$. Indeed, if we assume that $\|S(f)-\phi_f(j)\| > \varepsilon$, for all $j \in A$, then the quantum algorithm cannot have accuracy ε with probability at least $\frac{3}{4}$, and we reach a contradiction.

The triangle inequality yields that $\|S(f) - \phi_f(j)\| \leq \|S(f)-\phi_f(j^*)\| + \|\phi_f(j^*)-\phi_f(j)\| \leq 3\varepsilon$, for any $j \in A$. Hence, we have obtained a deterministic classical algorithm that solves the problem with error 3ε.

By our assumption, the classical components of the quantum algorithm may contain a number of function evaluations up to $n(\varepsilon)$ which implies

$$2n(\varepsilon) \geq \mathrm{comp}^{\mathrm{query}}_{\mathrm{clas}}(3\varepsilon). \qquad (9)$$

Since the quantum algorithm must have at least $\log_2 n(\varepsilon)$ qubits, as we indicated at the beginning of the proof, equation (9) implies that the qubit complexity of the quantum algorithm is bounded from below as follows

$$\mathrm{comp}^{\mathrm{qubit}}(\varepsilon) \geq \log_2 n(\varepsilon) \geq \log_2\left[\mathrm{comp}^{\mathrm{query}}_{\mathrm{clas}}(3\varepsilon)\right] - 1$$

and the proof is complete. ■

As we have already indicated quantum algorithms may have several measurements. They are sequences of quantum algorithms with a single measurement, i.e., a sequences of algorithms of the form (2), and the resulting algorithm has success probability, say, $\frac{3}{4}$. The individual quantum algorithms may use different numbers of qubits, and we denote by k the maximum of these numbers. One may reduce k not only at the expense of classical function evaluations but also by considering extremely long sequences of quantum algorithms with a single measurement. Therefore, to exclude such trivial cases we will assume that the total number of classical function evaluations used by the classical components of a sequence of quantum algorithms is a polynomial in 2^k, and so is the number of quantum algorithms with a single measurement that have been combined together to form the quantum algorithm with several measurements. Under these conditions we have the following corollary.

Corollary IV.1. *The qubit complexity of a quantum algorithm with several measurements is bounded as*

$$\mathrm{comp}^{\mathrm{qubit}}(\varepsilon) = \Omega\left(\log_2\left[\mathrm{comp}^{\mathrm{query}}_{\mathrm{clas}}(3\varepsilon)\right]\right).$$

[1] D. R. Simon, SIAM J. Comput. **26**, 1474 (1997).
[2] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, SIAM J. Computing **26(5)**, 1510 (1997).
[3] L. Grover, Phys. Rev. Lett. **79(2)**, 325 (1997), quant-ph/9706033.
[4] P. W. Shor, SIAM J. Comput. **26(5)**, 1484 (1997).
[5] J. F. Traub, G. W. Wasilkowski, and H. Woźniakowski, *Information-Based Complexity* (Academic Press, 1988).
[6] J. F. Traub and A. G. Werschulz, *Complexity and Information* (Cambridge University Press, 1998).
[7] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf, Proceedings FOCS'98 p. 352 (1998), quant-ph/9802049.
[8] E. Bernstein and U. Vazirani, SIAM J. Computing **26(5)**, 1411 (1997).
[9] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Phil. Trans. R. Soc. Lond. A. (1996).
[10] S. Heinrich, J. Complexity **18(1)**, 1 (2002), quant-ph/0105116.
[11] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
[12] S. Heinrich, J. Complexity **19**, 19 (2003).
[13] S. Heinrich, J. Complexity **20**, 5 (2004), quant-ph/0305030.
[14] S. Heinrich, J. Complexity **20**, 27 (2004), quant-ph/0305031.
[15] B. Z. Kacewicz, J. Complexity **21(5)**, 740 (2004).
[16] E. Novak, J. Complexity **17**, 2 (2001), quant-ph/0008124.
[17] A. Papageorgiou and H. Woźniakowski, Quantum Information Processing **4(2)**, 87 (2005), quant-ph/0502054.
[18] J. F. Traub and H. Woźniakowski, Quantum In-

formation Processing **1(5)**, 365 (2002), quant-ph/0109113.

[19] H. Woźniakowski, *The quantum setting with randomized queries for continuous problems* (2005), in progress.